What is Intelligence?

Something that characterizes humans and show rational behavior or best behavior is intelligence. It is ability to apply knowledge in order to perform better in an environment.

- The ability to reason
- The ability to understand
- The ability to create
- The ability to learn from experience
- The ability to plan and execute complex tasks

Human mind:

- Can use common sense and past experience for problem solving.
- Can understand natural language processing.
- Can percept environment like see, hear
- Can learn from experiences and acquire knowledge
- Can reason about facts and deduce new facts.

What is Artificial Intelligence?

Artificial Intelligence is a branch of Science which deals with helping machines to find solutions to complex problems in a more human-like fashion. This generally involves borrowing characteristics from human intelligence, and applying them as algorithms in a computer friendly way. A more or less flexible or efficient approach can be taken depending on the requirements established, which influences how artificial the intelligent behavior appears.

- "Giving machines ability to perform tasks normally associated with human intelligence."
- Branch of computer science that aims to create intelligence of machines
- Part of computer science concerned with designing intelligent machines
- Deals with issues like inference, reasoning, problem solving, knowledge representation, planning, natural language processing, perceptron, etc.

AI is generally associated with Computer Science, but it has many important links with other fields such as Math, Psychology, Cognition, Biology and Philosophy, among many others. Our ability to combine knowledge from all these fields will ultimately benefit our progress in the quest of creating an intelligent artificial being.

AI was coined by John McCathy in 1956 who defines AI as science and engineering of making intelligent machines especially intelligent computers.

Different definitions of AI are given by different books/writers. These definitions can be divided into two dimensions.

System that thinks like humans	**System that think rationally**
System that acts like humans	System that act rationally

Top dimension is concerned with thought processes and reasoning, where as bottom dimension addresses the behavior.

The definition on the left measures the success in terms of fidelity of human performance, whereas definitions on the right measure an ideal concept of intelligence, which is called rationality.

Human-centered approaches must be an empirical science, involving hypothesis and experimental confirmation. A rationalist approach involves a combination of mathematics and engineering.

Acting Humanly: The Turing Test Approach

- proposed by Alan Turing (1950)
- designed to convince the people that whether a particular machine can think or not
- a test based on indistinguishability from undeniably intelligent entities- human beings
- Involves an interrogator who interacts with one human and one machine. Within the given time the interrogator has to find out which one is human and which one is machine

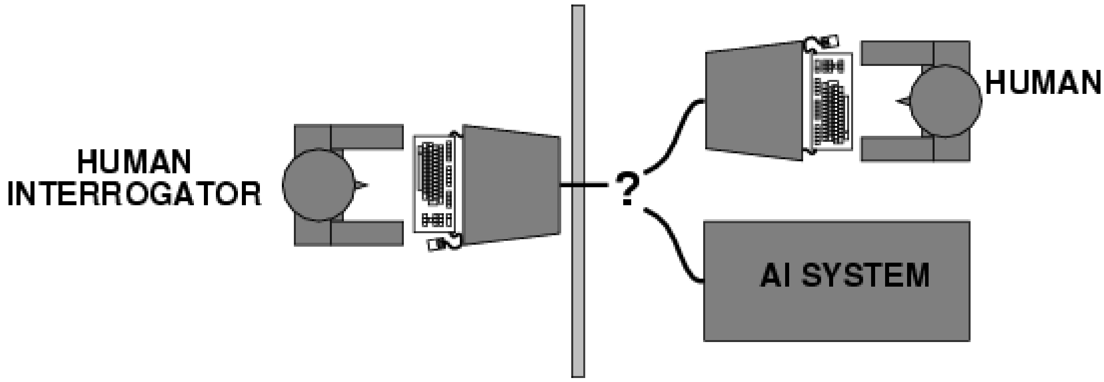

- The computer passes the test if a human interrogator can not tell whether the written response come from human or not
- To pass a Turing test, a computer must have following capabilities:

➢ Natural language processing

➢ Knowledge representation

➢ Automated reasoning

➢ Machine learning

Turing test avoid the physical interaction with human interrogator. Physical simulation of human beings is not necessary for testing the intelligence.

The total Turing test includes video signals and manipulation capability so that the interrogator can test the subject's perceptual abilities and object manipulation ability.
To pass the total Turing test computer must have following additional capabilities:

➢ Computer Vision: To perceive objects

➢ Robotics: To manipulate objects and move

Chinese room argument

- Searle (1999) summarized the Chinese Room argument
- Imagine a native English speaker who knows no Chinese locked in a room full of boxes of Chinese symbols (a data base) together with a book of instructions for manipulating the symbols (the program).
- Imagine that people outside the room send in other Chinese symbols which, unknown to the person in the room, are questions in Chinese (the input).
- And imagine that by following the instructions in the program the man in the room is able to pass out Chinese symbols which are correct answers to the questions (the output).
- The program enables the person in the room to pass the Turing Test for understanding Chinese but he does not understand a word of Chinese.

Thinking Humanly: Cognitive approach

Make the machines with mind. Cognition means the action or process of acquiring knowledge and understanding through thought, experience and senses. To make a machine that think like human brain, scientific theories of internal brain activities (cognitive model) are required. Two ways of doing this is:

– Predicting and testing human behavior (cognitive science)

– Identification from neurological data (Cognitive neuroscience)

Once we have precise theory of mind, it is possible to express the theory as a computer program. But unfortunately until up to now there is no precise theory about thinking process of human brain. Therefore it is not possible to make the machines that think like human brain

Thinking rationally: The laws of thought approach

Aristotle was one of the first who attempt to codify the right thinking that is irrefutable reasoning process. He gave Syllogisms that always yielded correct conclusion when correct premises are given.
For example:

Ram is a man
Man is mortal
->Ram is mortal
Let
p(x) -> x is man
q(x) -> x is mortal

then above statement can be written as

p(x) => q(x) Man is mortal
p(Ram) Ram is man

Then from modus ponens q(x) ia also true. That is Ram is mortal
This study initiated the field of logic. The logicist tradition in AI hopes to create intelligent systems using logic programming.

Problems:
It is not easy to take informal knowledge and state in the formal terms required by logical notation, particularly when knowledge is not 100% certain.

Solving problem principally is different from doing it in practice. Even problems with certain dozens of fact may exhaust the computational resources of any computer unless it has some guidance as which reasoning step to try first.

Acting Rationally: The rational Agent approach:

Agent is something that acts. Computer agent is expected to have following attributes:

- Autonomous control
- Perceiving their environment
- Persisting over a prolonged period of time
- Adapting to change
- And capable of taking on another's goal

Rational behavior means doing the right thing. The right thing is that which is expected to maximize goal achievement, given the available information. Rational Agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome. In this approach the emphasis is given to correct inferences.

One way to act rationally is to reason logically to the conclusion and act on that conclusion. On the other hand there are also some ways of acting rationally that can not be said to involve inference. For Example, recoiling from a host stove is a reflex action that is usually more successful than a slower action taken after careful deliberation.

Advantages: It is more general than laws of thought approach, because correct inference is just one of several mechanisms for achieving rationality.

It is more amenable/ controllable to scientific development than are approaches based on human behavior or human thought because the standard of rationality is clearly defined and completely general.

Applications of AI

Game playing
You can buy machines that can play master level chess for a few hundred dollars. There is some AI in them, but they play well against people mainly through brute force computation--looking at hundreds of thousands of positions. To beat a world champion by brute force and known reliable heuristics requires being able to look at 200 million positions per second.

Speech recognition
In the 1990s, computer speech recognition reached a practical level for limited purposes. Thus United Airlines has replaced its keyboard tree for flight information by a system using speech recognition of flight numbers and city names. It is quite convenient. On the the other hand, while it is possible to instruct some computers using speech, most users have gone back to the keyboard and the mouse as still more convenient.

Understanding natural language
Just getting a sequence of words into a computer is not enough. Parsing sentences is not enough either. The computer has to be provided with an understanding of the domain the text is about, and this is presently possible only for very limited domains.

Computer vision
The world is composed of three-dimensional objects, but the inputs to the human eye and computers' TV cameras are two dimensional. Some useful programs can work solely in two dimensions, but full computer vision requires partial three-dimensional information that is not just a set of two-dimensional views. At present there are only limited ways of representing three-dimensional information directly, and they are not as good as what humans evidently use.

Expert systems
A ``knowledge engineer'' interviews experts in a certain domain and tries to embody their knowledge in a computer program for carrying out some task. How well this works depends on whether the intellectual mechanisms required for the task are within the present state of AI. When this turned out not to be so, there were many disappointing results. One of the first expert systems was MYCIN in 1974, which diagnosed bacterial infections of the blood and suggested treatments. It did better than medical students or practicing doctors, provided its limitations were observed. Namely, its ontology included bacteria, symptoms, and treatments and did not include patients, doctors, hospitals, death, recovery, and events occurring in time. Its interactions depended on a single patient being considered. Since the experts consulted by the knowledge engineers knew about patients, doctors, death, recovery, etc., it is clear that the knowledge engineers forced what the experts told them into a predetermined framework. In the present state

of AI this has to be true. The usefulness of current expert systems depends on their users having common sense.

Heuristic classification
One of the most feasible kinds of expert system given the present knowledge of AI is to put some information in one of a fixed set of categories using several sources of information. An example is advising whether to accept a proposed credit card purchase. Information is available about the owner of the credit card, his record of payment and also about the item he is buying and about the establishment from which he is buying it (e.g., about whether there have been previous credit card frauds at this establishment).

Foundations of AI

Different fields have contributed to AI in the form of ideas, viewpoints and techniques.

Philosophy:
Logic, reasoning, mind as a physical system, foundations of learning, language and rationality.

Mathematics:
Formal representation and proof algorithms, computation, undecidability, intractability, probability.

Psychology:
adaptation, phenomena of perception and motor control.

Economics:
formal theory of rational decisions, game theory.

Linguistics:
Knowledge representation, grammar

Neuroscience:
Physical substrate for mental activities

Control theory:

Homeostatic systems, stability, optimal agent design

Sociology: Social behavior and social values, traditions

Omniscience and AI

The two concepts do not directly relate and it is not easy to compare and contrast.

Artificial intelligence is traditionally used by electronic devises to create an intelligent agent that can perceive and respond to data. For example a computer playing chess uses artificial intelligence (in the form of algorithms and specialized hardware) to determine the best move to make. Artificial intelligence can also be used by a computer to interpret information being taken in from a video camera to identify objects viewed.

Omniscience is the capacity to have unlimited knowledge about all things.

One logical conclusion of omniscience is that anything that is omniscient has no need to be intelligent.

Knowing everything means that for any possible situation or question, you already have the solution. There is no need to involve logic, rational thought, or such. An omniscient being need only be able to perceive, lookup the solution, and act on it.

In other words, it is only a big infinite database. Just as if you were to solve chess, all you'd need to do is encode each possible position, and when your turn comes up, record the positions of the pieces, and find the corresponding move to bring about the desired result.

Consider it solving reality.

Consider an AI that could understand us better than we understand ourselves. Possessing such a deep understanding of how we think, work, and our skills and abilities to such an extent where it could help us make decisions that would benefit us in the long run. It could direct our lives and we could follow its recommendations and believe that it knows better than we do. The algorithms and amount of data it processes creating a science of success, a systematic plan for life satisfaction.

At first blush, that seems quite powerful and enticing. Why is that idea so attractive? It would be nice if we could take some of the uncertainty out of life. Make tough decisions easier with predictions based on processing data we do not fully understand. When we make decisions we decide based on the information available to us, consider as many implications as possible, but if we were able to know the best route instantly based on a hundred, a thousand or a million times more variables than we could conceive of or even have access to, would not that result in a more trustworthy result? It seems like it would be quite convenient to have a super intelligent AI life coach.

If it said, “Don’t do drugs.” or “You should stay in and read Aristotle’s Ethics tonight rather than go to that party.” or “Put in 12 hours of overtime this weekend to get the project done on time.” or “You should buy the car off Craig’s list rather than lease the Tesla.” or “Put the money in savings rather than buy the _____ you were looking at on Amazon.” Maybe even, “Don’t divorce your spouse, work it out.” Would we even listen to its advice?

Someone would have to program it, design the algorithms, consider how to weigh outcomes and variables. Could we trust that person? What if their values were different from ours? No matter how many variables you examine, it cannot be omniscient and know the future. There is just no way to account for everything.

Brief history of AI

What happened after WWII?
– 1943: Warren Mc Culloch and Walter Pitts: a model of artificial boolean neurons to perform computations.
– First steps toward connectionist computation and learning (Hebbian learning).
– Marvin Minsky and Dann Edmonds (1951) constructed the first neural network computer
– 1950: Alan Turing's "Computing Machinery and Intelligence"
– First complete vision of AI.

The birth of AI (1956)
- Dartmouth Workshop bringing together top minds on automata theory, neural nets and the study of intelligence.
– Allen Newell and Herbert Simon: The logic theorist (first nonnumeric thinking program used for theorem proving)
– For the next 20 years the field was dominated by these participants.

Great expectations (1952-1969)
– Newell and Simon introduced the General Problem Solver.
– Imitation of human problem-solving
– Arthur Samuel (1952-)investigated game playing (checkers) with great success.
– John McCarthy(1958-) :
– Inventor of Lisp (second-oldest high-level language)
– Logic oriented, Advice Taker (separation between knowledge and reasoning)

– Marvin Minsky (1958 -)
– Introduction of microworlds that appear to require intelligence to solve: e.g. blocks-world.
– Anti-logic orientation, society of the mind.

Collapse in AI research (1966 - 1973)
– Progress was slower than expected.
– Unrealistic predictions.
– Some systems lacked scalability.
– Combinatorial explosion in search.
– Fundamental limitations on techniques and representations.
– Minsky and Papert (1969) Perceptrons.

AI revival through knowledge-based systems (1969-1970)
– General-purpose vs. domain specific

- E.g. the DENDRAL project (Buchanan et al. 1969)
First successful knowledge intensive system.

– Expert systems
- MYCIN to diagnose blood infections (Feigenbaum et al.)
- Introduction of uncertainty in reasoning.
– Increase in knowledge representation research.

- Logic, frames, semantic nets, …

AI becomes an industry (1980 - present)
– R1 at DEC (McDermott, 1982)
– Fifth generation project in Japan (1981)
– American response …

Puts an end to the AI winter.

Connectionist revival (1986 - present): (Return of Neural Network)
– Parallel distributed processing (RumelHart and McClelland, 1986); backprop.

AI becomes a science (1987 - present)
– In speech recognition: hidden markov models
– In neural networks
– In uncertain reasoning and expert systems: Bayesian network formalism

The emergence of intelligent agents (1995 - present)
– The whole agent problem: "How does an agent act/behave embedded in real environments with continuous sensory inputs"